A PARALLEL 3D DIGITAL WAVEGUIDE MESH MODEL WITH TETRAHEDRAL TOPOLOGY FOR ROOM ACOUSTIC SIMULATION

Guilherme Campos and David Howard

Music Technology Research Group
Department of Electronics - University of York
agrc101@york.ac.uk

ABSTRACT

Following a summary of the basic principles of 3D waveguide mesh modelling and the context of its application to room acoustic simulation, this paper presents a detailed analysis of the tetrahedral mesh topology and describes its implementation on a parallel computer model. Its structural characteristics are analysed, with particular emphasis on how they influence execution speed. Performance deterioration due to communication overhead in the parallelised model is discussed.
Theoretical predictions are compared with data from performance tests carried out on different computer platforms and both are contrasted with the corresponding results from the rectilinear model, in order to assess the practical efficiency of the model.
Objective validation tests are reported and discussed.

1. INTRODUCTION

The applications of room acoustic modelling range from auditoria design to the development of musical equipment and the creation of tools for sound synthesis and manipulation. Although varying degrees of accuracy may be required by different applications, *auralisation* can be regarded as the ultimate objective.

The problem of auralisation amounts to solving the sound wave equation for a particular room under analysis. Unfortunately, analytical solutions can be obtained only for very simple geometries and idealised boundary conditions [1].

Models based on numerical methods, known as *physical models*, have to be used. Excessive computational loading has been their main drawback for practical application in room acoustic modelling.

Simplified approaches to the analysis of sound propagation have been proposed to circumvent this problem, namely *geometric models* inspired by optics [1]. However, their results are far from satisfactory when a high level of perceptual accuracy is required [2], as is the case in the acoustical reconstruction of ancient buildings or structures for musical or musicological purposes.

Finding ways of making physical models practical for room acoustic simulation is therefore a research priority. The work reported in this paper is based on *digital waveguide modelling*, a Finite Difference Time Domain (FDTD) method developed for musical applications by Van Duyne and Smith [3]. Being a physical modelling technique, it automatically accounts for all wave propagation phenomena, including diffraction.

2. DIGITAL WAVEGUIDE MODELLING

By discretising time and space, the *travelling wave solution* to the 1-D wave equation for either flow or pressure can be implemented digitally with a bi-directional pair of delay lines. Such a structure is called a *digital waveguide* [3].

The point of intersection of n digital waveguides is called a *scattering junction* or *node*. Using scattering junctions, multi-dimensional digital *waveguide meshes* can be constructed [3]. In this paper, superscripted + and – denote wave components travelling respectively to and from a junction.

Assuming lossless transmission (i.e. neglecting energy absorption by the propagation medium), the sound pressure (p) at a scattering junction can be expressed as a function of incoming sound pressure travelling waves (p_i^+) [4]:

$$p = \frac{2\sum_{i=1}^{n} p_i^+ / Z_i}{\sum_{i=1}^{n} 1/Z_i} \qquad (1)$$

where Z_i is the acoustic impedance along each line of propagation.

The sound pressures ($p_i = p_i^+ + p_i^-$) in all crossing waveguides are equal at the junction [4]:

$$p_1 = p_2 = ... = p_n = p \qquad (2)$$

In a *regular* mesh, all waveguide segments between nodes have the same length, being called *delay units*. Wave components take one *time step T* to travel the distance through a delay unit from a node i to the *opposite* one i,opp [4]:

$$p_i^+[nT] = p_{i,opp}^-[(n-1)T] \qquad (n \in N) \qquad (3)$$

Strings and air columns can be modelled using one-dimensional *digital waveguides*. Higher-dimensional models, known as *waveguide meshes*, can be formed by digital waveguides interconnected in regular arrangements [3].

The application of two-dimensional models to the simulation of acoustic membranes and percussion instruments has been particularly successful. Very rich and natural-sounding timbres can be obtained [5]. This adds to the idea that 3D meshes have the potential to provide accurate room simulation.

The main handicap of the model, common to any FDTD method, is *dispersion error*. This means that wave propagation speed is variable with frequency and direction of propagation.

Higher-frequency components generally lag behind, causing direction-dependent signal distortion [3].

This error, as well as those caused by finite spatial resolution in the representation of boundaries, can be reduced by increasing the mesh density, but at the expense of computation time

Several 2D mesh topologies have been studied, particularly *rectilinear*, *hexagonal* and *triangular*, the aim being to find the best compromise between dispersion error, computation speed and ease of implementation [5].

Similar efforts are being applied to the 3D case. This paper explores the *tetrahedral* structure proposed by the originators of the waveguide-mesh modelling technique [6]. They show that tetrahedral models compute a valid finite difference approximation to the 3D wave equation and point out several potential advantages which could lead to practical application in room acoustic simulation.

3. THE TETRAHEDRAL MESH STRUCTURE

Tetrahedral meshes replicate the molecular structure of the diamond crystal, with nodes corresponding to carbon nuclei and the tetrahedrally-spaced bi-directional delay units around each node corresponding to the chemical bonds with neighbouring nuclei [6]. Figure 1 presents the geometry of the tetrahedral arrangement: the four neighbours of a node are vertices of a cube whose geometrical centre is the given node. The angle α between bonds is $2\arctan\sqrt{2} \cong 109.47°$. The inter-nodal distance is $d=\sqrt{3}\,s/2$, s being the edge length of the tetrahedral *site* [7].

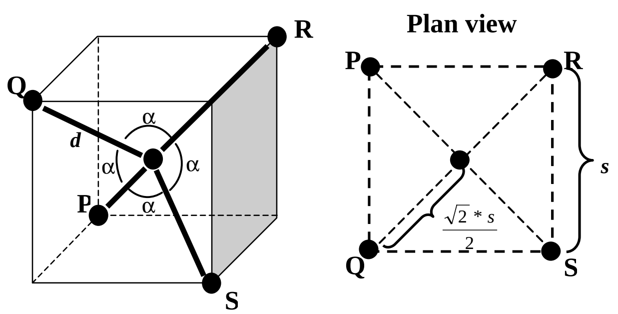

Figure 1. *A tetrahedral site*

An important concept in the analysis of crystalline structures is that of the *unit cell – the smallest repeating unit which shows the full symmetry of the crystal structure* [8].

The unit cell of the diamond crystal, a *cubic* structure, is represented in figure 2*a*. It is formed by 8 identical cubes, 4 of them occupied by tetrahedral sites, the remainder being empty.

Of the 18 nodes shown, 4 are in the interior of the cube, 6 at the faces (shared by 2 unit cells) and 8 at the corners (shared by 8 unit cells). Therefore, a unit cell comprises 4+(1/2).6+(1/8).8 = 8 nodes [9]. The same observation can be made considering a small displacement of the unit cell's cubic contour by a vector (ε, ε, ε), with $|\varepsilon| < s/2$, all nodes remaining fixed. Only 8 nodes (one at a vertex, three on the faces forming it and of course the four interior nodes) remain inside the cube. If $\varepsilon < 0$, they correspond to the points marked A to H. $|\varepsilon| = s/4$ gives the most uniform node distribution within the cell, as shown in figure 2*b*.

Every delay unit is aligned with one of the model's four propagation axes:

1. $x = y = z$ 2. $x = y = -z$ 3. $x = -y = z$ 4. $-x = y = z$

Figure 2*a* shows there are 16 bi-directional delay units per unit cell or, equivalently, 2 per node. Figure 2*b* highlights the 7 bonds between nodes within the same unit cell.

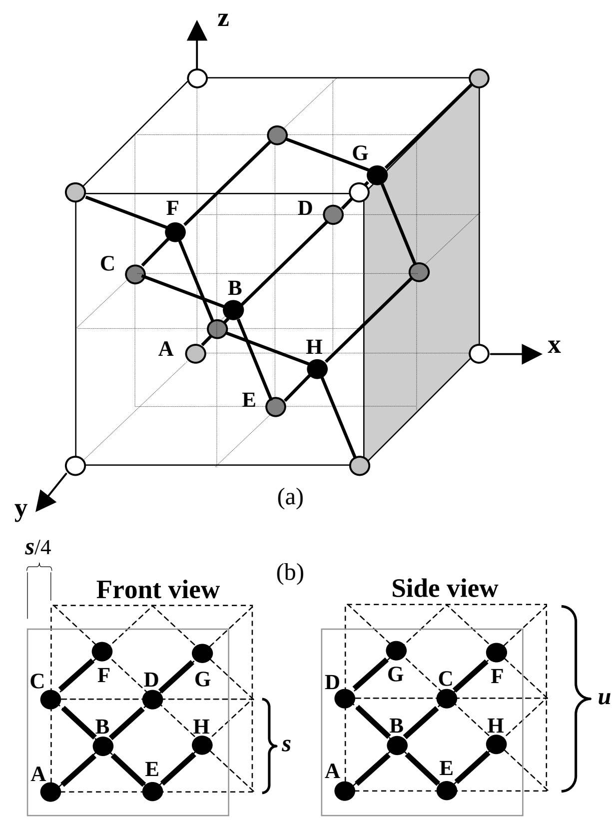

Figure 2. *A unit cell of the tetrahedral mesh* [9]

A mesh is formed by unit cells placed adjacent to one another, as figure 3 illustrates.

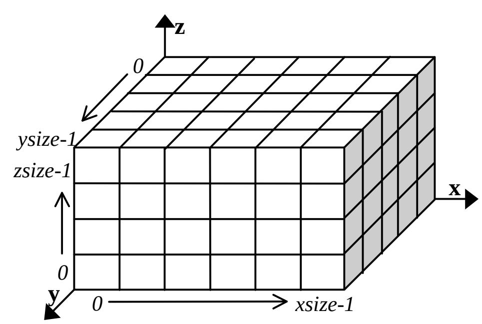

Figure 3. *A tetrahedral mesh formed by a 3D [xsize.ysize.zsize] array of unit cells*

Considering only an infinitesimal cell boundary contour displacement, $\lim_{\varepsilon\to 0^-}$ (ε, ε, ε), fig. 2*a* holds and the node coordinates of a generic unit cell (i, j, k) in the model of fig. 3 can be worked out as:

A $(i, j, k) \rightarrow 2s(i, j, k)+(0, 0, 0)$;
B $(i, j, k) \rightarrow 2s(i, j, k)+ (s/2, s/2, s/2)$;
C $(i, j, k) \rightarrow 2s(i, j, k)+(0, s, s)$;
D $(i, j, k) \rightarrow 2s(i, j, k)+ (s, 0, s)$;
E $(i, j, k) \rightarrow 2s(i, j, k)+ (s, s, 0)$;
F $(i, j, k) \rightarrow 2s(i, j, k)+ (s/2, 3s/2, 3s/2)$;
G $(i, j, k) \rightarrow 2s(i, j, k)+ (3s/2, s/2, 3s/2)$;
H $(i, j, k) \rightarrow 2s(i, j, k)+ (3s/2, 3s/2, s/2)$;

where $0 \leq i \leq xsize\text{-}1$, $0 \leq j \leq ysize\text{-}1$, $0 \leq k \leq zsize\text{-}1$ and $2s = u$ is the edge of the unit cell.

4. THE TETRAHEDRAL MODEL ALGORITHM

Each tetrahedral mesh node is represented by 10 fields, as detailed in table 1. Eight of them implement bi-directional communication ports with the 4 neighbouring nodes.

Field		Symbol	Data type
Node pressure (air nodes) / / Reflection factor (boundary nodes)		p / α	4-byte float
Propagation axis 1	Output port	p_1^+	4-byte float
	Input port	p_1^-	4-byte float
Propagation axis 2	Output port	p_2^+	4-byte float
	Input port	p_2^-	4-byte float
Propagation axis 3	Output port	p_3^+	4-byte float
	Input port	p_3^-	4-byte float
Propagation axis 4	Output port	p_4^+	4-byte float
	Input port	p_4^-	4-byte float
Configuration		node_config	char

Table 1. *Node structure*

As in the 3D rectilinear model [10], nodes can be configured (*node_config*) as *air nodes* or *boundary nodes*. The field that holds the wave variable p in the former case, is used to store the reflection factor α of the corresponding surface in the latter. *Air* nodes can be assigned special functions, namely *mesh excitation* (modelling sound *sources*) and *output* (modelling sound *receivers*). This allows the acoustic characteristics of any room with various surface coverings and different source and receiver locations to be modelled.

The modelling algorithm is the iterative two-pass computation for lossless propagation [3] applied before [10] to the 3D rectilinear topology.

The **scattering pass** calculations for boundary nodes remain exactly the same, but for air nodes equations 1 and 2 yield:

$$p = \frac{1}{2}(p_1^+ + p_2^+ + p_3^+ + p_4^+) \tag{4}$$

$$p_i^- = p - p_i^+ \qquad i \in \{1, 2, 3, 4\} \tag{5}$$

The **delay pass** implements equation 3 for every delay unit in the mesh. Data transfer is more complex than in the rectilinear mesh, because of the eight different node positions in a unit cell. Table 2 identifies the neighbours corresponding to each position.

	Nodes in cell (i, j, k)							
	A	**B**	**C**	**D**	**E**	**F**	**G**	**H**
i, j, k	B	A/E/D/C	F/B	G/B	H/B	C	D	E
i, j, k-1					F/G			
i, j, k+1						E	E	
i, j-1, k				F/H				
i, j+1, k						D		D
i-1, j, k			G/H					
i+1, j, k							C	C
i, j-1, k-1	F							
i-1, j, k-1	G							
i-1, j-1, k	H							
i, j+1,k+1						A		
i+1, j, k+1							A	
i+1, j+1, k								A

Table 2. *Map of node interconnections.*

There are links not only to the 6 unit cells with faces adjacent to it (as was the case in the rectilinear model) but also to 6 others positioned 'diagonally' i.e. sharing only one edge with it, as figure 4 illustrates.

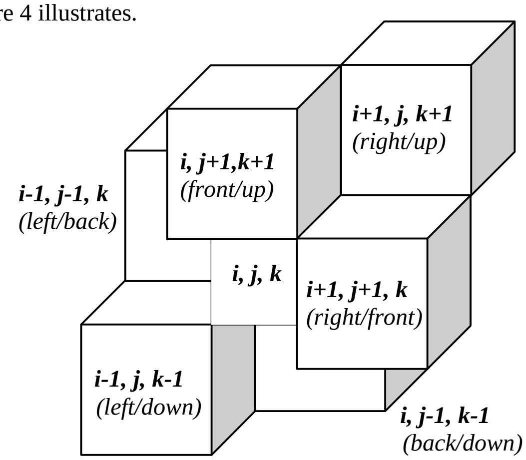

Figure 4. *Relative position of the six neighbouring unit cells exchanging data with a generic unit cell (i, j, k), through a common edge*

5. PARALLEL IMPLEMENTATION

A computer model based on the tetrahedral structure described above was implemented using the PVM (Parallel Virtual Machine) software package [11], allowing it to run indifferently on multiprocessors or workstation clusters.

The parallelisation strategy was exactly the same as described in [10] for a rectilinear mesh model. The overall model of the room is partitioned into a set of cuboid-shaped *model blocks* with identical dimensions.

A tetrahedral model block consists of a 3D array of unit cells, as represented in figure 3. Considering in turn all the nodes having the same designation (one per unit cell), it can be described as a set of eight interlaced 3D rectilinear [*xsize.ysize.zsize*] node arrays, corresponding to designations A through H; this was the implementation approach adopted.

Based on a map of their positions on a 3D grid, a *master* program spawns one *slave* task per block and enables it to communicate directly with the tasks modelling neighbouring blocks, for data transfer between surface nodes. This requires that model blocks be associated with appropriate sets of *communication buffers*.

In order to help improve performance, the computation sequence in each iteration cycle (fig. 5) is organised so that data required by neighbours is made available as early as possible and data required from them can be received as late as possible [10].

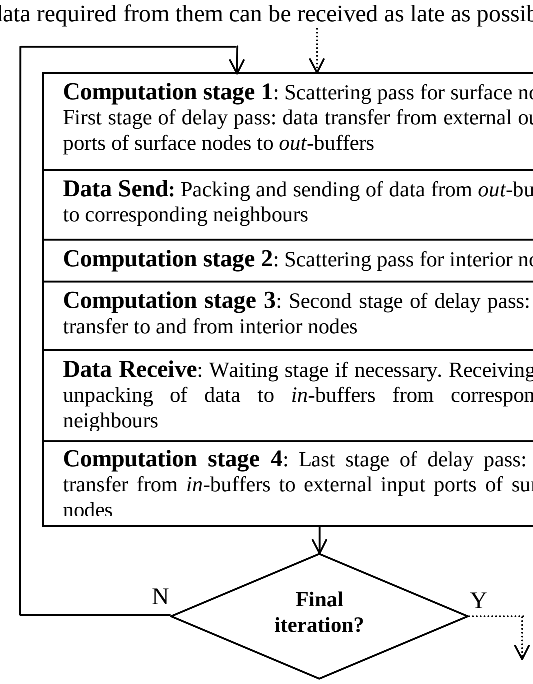

Figure 5. *Sequence of main computation stages in one cycle of parallelised operation.*[10]

The main difference to the rectilinear topology in terms of data transfer is that it occurs not only between any two blocks with a common face but also between blocks having a common edge on certain positions. The symmetry properties exhibited at unit cell level are reflected on a larger scale – in effect, figure 4 is applicable to inter-block data transfer, simply replacing the caption words *unit cell(s)* by *model block(s)*.

The fact that a generic tetrahedral model block is composed of 8 arrays and exchanges data with 12 neighbours, as opposed to only one array and 6 active neighbours for its rectilinear counterpart, gives rise to a relatively large number of different inter-block communication situations. As a result, a total of 30 input/output pairs of communication buffers are necessary, in contrast with a mere 6 in a rectilinear model block.

6. COMPUTATION TIME

Van Duyne and Smith [6] show that, as in rectilinear meshes, the apparent wave speed is $\sqrt{1/3}$ space units per time sample (the space unit being the internodal distance, d):

$$d = \frac{\sqrt{3}c}{f_u} \tag{6}$$

where c represents the speed of sound and $f_u = 1/T$ is the *update frequency* of the mesh, T being the time step between consecutive iterations. As f_u defines the rate of injection of input samples, it is equivalent to the audio sampling rate, f_s.

Dispersion error [6], inherent to FDTD methods, means that propagation characteristics are direction-dependent; the only directions in which all frequencies travel at the same correct speed are now those of the coordinate axes. For these, equation 6 can be confirmed by observing that waves take four time steps ($4T$) to traverse the edge of the unit cell ($4d/\sqrt{3}$).

According to equation 6, for the same sampling frequency, tetrahedral and rectilinear models of a given room have the same inter-nodal distance. This results in rectilinear meshes being denser. In fact, in a tetrahedral mesh, as there are eight nodes per unit cell, the volume per node, v, is that of a tetrahedral site:

$$v_{tetrahedral} = s^3 = \left(\frac{2d}{\sqrt{3}}\right)^3 \tag{7}$$

whereas in a rectilinear mesh it is given by:

$$v_{rectilinear} = d^3 \tag{8}$$

Mesh density (nodes per volume unit) can be measured by $1/v$. The density ratio between the tetrahedral and rectilinear meshes for the same inter-nodal distance can be calculated as:

$$\frac{(1/v)_{tetrahedral}}{(1/v)_{rectilinear}} = \frac{\left(\frac{\sqrt{3}}{2d}\right)^3}{\left(\frac{1}{d}\right)^3} = \left(\frac{\sqrt{3}}{2}\right)^3 \cong 65\% \tag{9}$$

The condition for a rectilinear block with dimensions r_x, r_y, r_z (nodes) to model the same volume as a tetrahedral one with dimensions t_x, t_y, t_z (unit cells) at a given sampling rate is:

$$r_i d = \frac{4d}{\sqrt{3}} t_i \Leftrightarrow r_i = \frac{4}{\sqrt{3}} t_i \qquad i \in \{x, y, z\} \tag{10}$$

Only approximate solutions can be obtained, as r_i and t_i are integers; the corresponding blocks are considered *equivalent*.

A formula for the computation time of a 3D rectilinear model is derived in [10]. In a similar way, using equations 6 and 7, the corresponding formula for the tetrahedral mesh can be obtained:

$$T_M = \left(\frac{44100}{2c}\right)^3 .V .RT_{60} . f_n^4 .t_n \tag{11}$$

where V is the room volume (m^3), RT_{60} is a measure of reverberation time (s), t_n=*44100t* is the average computation time per node per 44100 iterations (s), corresponding to 1s of audio at the reference frequency, and f_n is a 'normalised' sampling rate (adimensional) given by:

$$f_n = \frac{f_S}{44100Hz} \tag{12}$$

The only difference between the two computation time formulae is the numerical factor, 65% smaller in the tetrahedral mesh as a direct result of its lower density.

In addition, t_n should be lower:

1) The scattering pass (equations 4 and 5) involves only 7 algebraic additions and one division by 2, as opposed to 11 additions and a division by 3 with the rectilinear model 6-port junctions. The division by 2 cannot be implemented here through bit-shifting as calculations are all floating-point.

2) As there are only 2 delay units per node (3 in the rectilinear mesh), the number of delay pass operations also decreases by 1/3.

Thus, with t_n reduced by approximately 1/3 (see figure 5), T_M would be 57.1% lower than in the rectilinear mesh.

In parallelised operation, the computation time is affected by inter-block data transfer. The ratio *S/V* between the number of surface delay units (connecting to exterior nodes) and the total number of nodes gives a clear indication of *communication overhead*. This ratio, a decreasing function of block size, is minimised for cubic shapes in both the rectilinear and tetrahedral topologies. For equivalent cubic blocks, it can be shown that:

$$\frac{(S/V)_{tetrahedral}}{(S/V)_{rectilinear}} = \frac{2}{\sqrt{3}}\left(1 - \frac{1}{\sqrt{3}n}\right) \cong \frac{2}{\sqrt{3}} \qquad (13)$$

where n is the edge of the rectilinear block, in number of nodes. This indicates a higher communication overhead in the tetrahedral mesh. The larger number of active neighbours and communication buffers (respectively 12 and 60 against only 6 and 12 in the rectilinear case) is likely to aggravate the situation, especially for small blocks.

7. PERFORMANCE TESTS

In order to assess the efficiency of the model and the effect of communication overhead, the same technique as used in [10] for the rectilinear model was adopted. Two simple tests were carried out on identical *SGI O2* workstations connected by a network (100 Mbit/s Ethernet). In the first test, only one cubic block (no neighbours) was tested under PVM on one of the machines, with all send-receive operations disabled. In the second test, 2 tasks were spawned, one on each workstation, both modelling blocks identical to the previous one. Blocks were made to interact as if both had the maximum possible number of active neighbours (12), by sending all the corresponding buffers to each other, therefore simulating a worst-case scenario in terms of communication overhead.

To allow a comparative analysis, the two tests were applied to sets of equivalent tetrahedral and rectilinear blocks with increasing size. The results are presented in figure 6.

The abscissa unit is scaled by f_n (see equation 12) to make the graph applicable to any sampling frequency (this also applies to figure 7).

Test 1 reveals a performance improvement in excess of the theoretical prediction of 33%, caused by delay pass calculations being faster than expected. This may be related to more efficient memory access, resulting from data being transferred between interlaced small-size arrays rather than within a single large array.

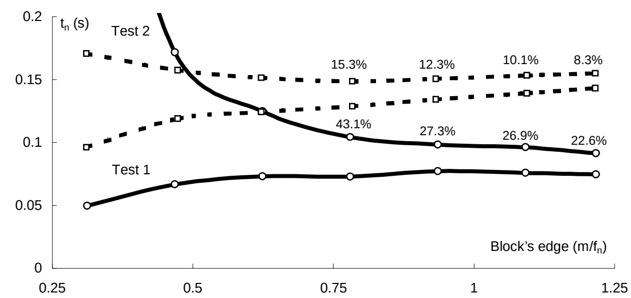

Figure 6. *Comparative communication overhead tests. The percentage values indicate performance deterioration in test 2 (worst-case inter-block communication) relative to the t_n values from test 1 (no inter-block communication)*

As predicted, after a steep decrease in the initial part of the graph, the communication overhead ratio between equivalent tetrahedral and rectilinear blocks stabilises at a value only slightly greater than that given by equation 13. The higher percentages seen for the tetrahedral mesh are due mainly to communication overhead being compared to much lower values of t_n.

Normal parallel model operation was tested on a much faster platform – a 4-node, 8-processor *SGI Origin 2000*. A set of 27 room models formed by 1, 2 and 4 identical tetrahedral cubic blocks of 9 increasing sizes were successively tested, with no more than one block running on a given processor.

The results are presented in figure 7 where they can be compared with those obtained with the rectilinear mesh for a set of equivalent rooms (see equation 10).

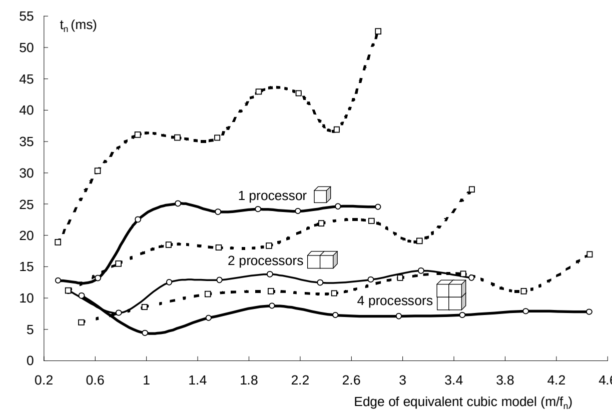

Figure 7. *Multiprocessor comparative performance of the tetrahedral mesh (solid lines) and rectilinear mesh (dashed lines).*

In order to try and avoid any possible bias, as multi-user access can cause random fluctuations in processor loading, the results were averaged over a large number of non-consecutive measurements, rectilinear and tetrahedral model runs being interleaved.

A very significant performance improvement can be observed for both topologies each time the number of processors is doubled, except for small models, when the effect of communication overhead is much more apparent than in the rectilinear mesh, increasing with number of blocks, as expected.

Performance gains relative to the rectilinear model are variable with block size, reaching more than 100% for the largest sizes tested.

8. MODEL VALIDATION

Using a model parallelised over 6 processors on the computer tested before, a 2*s* impulse response at 44.1*kHz* was obtained for a rectangular room with dimensions as indicated in figure 8. The source and the receiver were placed at opposite corners of the room model.

Figure 8 shows the initial portion of the frequency spectrum obtained by applying FFT to the last 1.8*s* of the impulse response, using uniform windowing for best frequency resolution.

There is an excellent correspondence between the modal frequencies calculated analytically and the spectral peaks, providing objective validation of the model. The slight differences observed, more apparent for higher frequencies, can be attributed to dispersion error [12].

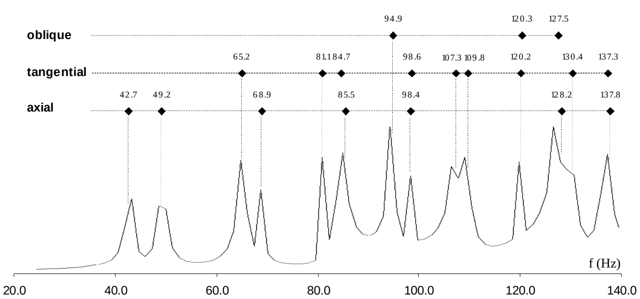

Figure 8. *Comparison between the spectral modes calculated analytically for a room with* W=4.025m, D=3.495m; H=2.496m *and the frequency response of the corresponding 3D tetrahedral model*

Preliminary listening tests with anechoically recorded sounds suggest that the model is behaving appropriately. In addition, binaural room responses obtained with a coarse model of a listener's head have confirmed that it is able to provide appropriate sound localisation cues.

9. CONCLUSIONS AND FURTHER WORK

A parallel computer implementation of the tetrahedral 3D waveguide mesh model for room acoustic simulation is described. Significant memory savings have been achieved and tests show that its computation time can be less than 3 times that of its rectilinear counterpart for the same computing platform, audio sampling rate and room volume. This performance improvement exceeds the theoretical predictions; this may result from a more efficient coding strategy. This aspect will be investigated in the future.

A study of the perceptual effects of spatial discretisation error in boundary representation is necessary in order to evaluate the impact of a decreased mesh density.

Large-scale parallelisation using the data decomposition strategy presented here requires more complex inter-block communication interfaces than with the rectilinear mesh as the number of active neighbours is doubled.

On the other hand, as communication overhead is higher and computational speed is faster a larger block size would be optimal.

10. REFERENCES

[1] Kuttruff, H. (1973) *Room Acoustics*. London, Applied Science Publishers.

[2] Murphy, D. T. and Howard, D. M. (1998) 'A multi-channel spatial simulation system for computer music applications' (*Proceedings of the Institute of Acoustics*, vol. 20, part 5, pp 161-168).

[3] Van Duyne, S. A. and Smith III, J. O. (1993) 'Physical Modeling with the 2-D Digital Waveguide Mesh' (*ICMC Proceedings*, pp 40-47).

[4] Savioja, L.; Karjalainen, M. and Takala, T. (1996) 'DSP Formulation of a Finite Difference Method for Room Acoustics Simulation' (*Proceedings of the NORSIG'96 – IEEE Nordic Signal Processing Symposium*, Espoo, Finland, 24-27 Sept., pp 455-458).

[5] Fontana, F. and Rocchesso, D. (1998) 'Physical Modelling of Membranes for Percussion Instruments' (*Acustica – acta acustica*, vol. 84, pp 529-542).

[6] Van Duyne, S. A. and Smith III, J. O. (1996) 'The 3-D Tetrahedral Digital Waveguide Mesh with Musical Applications' (*ICMC Proceedings,* pp 9-16).

[7] Adams, D. M. (1974) *Inorganic Solids – an Introduction to Concepts in Solid-state Structural Chemistry*. John Wiley & Sons.

[8] West, A. R. (1988) *Basic Solid State Chemistry*. John Wiley & Sons.

[9] Smart, L. and Moore, E. (1995) *Solid State Chemistry – an Introduction*. Chapman & Hall.

[10] Campos, G. and Howard, D. M. (2000) 'On the Computation Time of Three-Dimensional Digital Waveguide Mesh Acoustic Models' (*Proceedings of the 26th Euromicro Conference*, Maastricht, Holland, 5-7 Sept., vol. II, pp 332-339).

[11] Geist, A.; Beguelin, A.; Dongarra, J.; Jiang, W.; Manchek, R. and Sunderam, V. (1994) *PVM: Parallel Virtual Machine – A User's Guide and Tutorial for Networked Parallel Computing*. Cambridge, Massachussets, MIT Press.

[12] Savioja, L.; Backman, J.; Järvinen, A. and Takala, T. (1995) 'Waveguide Mesh Method for Low-Frequency Simulation of Room Acoustics' (*Proceedings of the 15th ICA (ICA 95)*, Trondheim, Norway, 26-30 June, vol. 2, pp 637-640).

Acknowledgement: The first author of this paper is sponsored by the Portuguese Government through FCT (programme PRAXIS XXI).